



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

| | | |
|---|-----------|--|
| (51) International Patent Classification ⁶ : G06K 9/00 | A1 | (11) International Publication Number: WO 00/05677 (43) International Publication Date: 3 February 2000 (03.02.00) |
|---|-----------|--|

(21) International Application Number: PCT/US98/15350

(22) International Filing Date: 23 July 1998 (23.07.98)

(71) Applicant: LOCKHEED MARTIN CORPORATION [US/US];
6801 Rockledge Drive, Bethesda, MD 20817-1836 (US).

(72) Inventors: SHAPIRO, Gary, Lee; 31232 Cree Drive, Evergreen, CO 80439 (US). OLIVER, Davis, ReGoalt, Jr.; 7621 South Bemis Street, Littleton, CO 80120 (US).

(74) Agent: CARSON, W., Scott; Dorr, Carson, Sloan & Bimey, P.C., 3100 East 6th Avenue, Denver, CO 80206 (US).

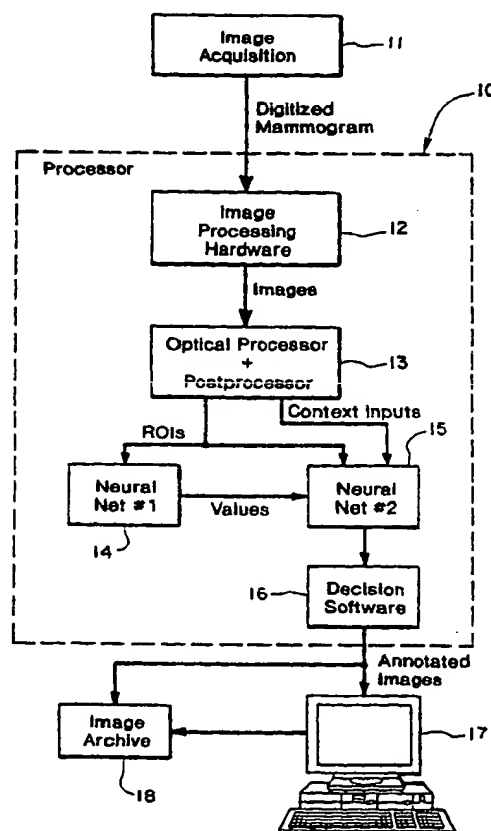
(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published*With international search report.*

(54) Title: SYSTEM FOR AUTOMATED DETECTION OF CANCEROUS MASSES IN MAMMOGRAMS

(57) Abstract

A system (10) for automated detection of cancerous masses in mammograms initially identifies regions of interest (ROIs) using Fourier analysis (e.g., by means of an optical correlator). Context data is extracted from the mammogram for each ROI, such as size, location, ranking, brightness, density, and relative isolation from other ROIs. The pixels in the ROI are averaged together to create a smaller array of super-pixels (24), which are input into a first neural net (14). A second neural net (15) receives the output values from the first neural net and the context data as inputs and generates an output score indicating whether the ROI contains cancerous mass. The second neural net can also be provided with context data from another view of the same breast, the same view of the other breast, or a previous mammogram for the same patient.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

| | | | | | | | |
|----|--------------------------|----|---------------------|----|-----------------------|----|--------------------------|
| AL | Albania | ES | Spain | LS | Lesotho | SI | Slovenia |
| AM | Armenia | FI | Finland | LT | Lithuania | SK | Slovakia |
| AT | Austria | FR | France | LU | Luxembourg | SN | Senegal |
| AU | Australia | GA | Gabon | LV | Latvia | SZ | Swaziland |
| AZ | Azerbaijan | GB | United Kingdom | MC | Monaco | TD | Chad |
| BA | Bosnia and Herzegovina | GE | Georgia | MD | Republic of Moldova | TG | Togo |
| BB | Barbados | GH | Ghana | MG | Madagascar | TJ | Tajikistan |
| BE | Belgium | GN | Guinea | MK | The former Yugoslav | TM | Turkmenistan |
| BF | Burkina Faso | GR | Greece | | Republic of Macedonia | TR | Turkey |
| BG | Bulgaria | HU | Hungary | ML | Mali | TT | Trinidad and Tobago |
| BJ | Benin | IE | Ireland | MN | Mongolia | UA | Ukraine |
| BR | Brazil | IL | Israel | MR | Mauritania | UG | Uganda |
| BY | Belarus | IS | Iceland | MW | Malawi | US | United States of America |
| CA | Canada | IT | Italy | MX | Mexico | UZ | Uzbekistan |
| CF | Central African Republic | JP | Japan | NE | Niger | VN | Viet Nam |
| CG | Congo | KE | Kenya | NL | Netherlands | YU | Yugoslavia |
| CH | Switzerland | KG | Kyrgyzstan | NO | Norway | ZW | Zimbabwe |
| CI | Côte d'Ivoire | KP | Democratic People's | NZ | New Zealand | | |
| CM | Cameroon | | Republic of Korea | PL | Poland | | |
| CN | China | KR | Republic of Korea | PT | Portugal | | |
| CU | Cuba | KZ | Kazakstan | RO | Romania | | |
| CZ | Czech Republic | LC | Saint Lucia | RU | Russian Federation | | |
| DE | Germany | LI | Liechtenstein | SD | Sudan | | |
| DK | Denmark | LK | Sri Lanka | SE | Sweden | | |
| EE | Estonia | LR | Liberia | SG | Singapore | | |

SYSTEM FOR AUTOMATED DETECTION OF CANCEROUS MASSES IN MAMMOGRAMS

5

BACKGROUND OF THE INVENTION

1. **Field of the Invention.** The present invention relates generally to the field of computer-aided diagnosis of medical images. More specifically, the present invention discloses an automated system for detecting cancerous masses in digital mammograms.

10

2. **Statement of the Problem.** While the problem of reducing breast cancer mortality is substantial, mammography provides an important tool for early detection. False negative rates are high, however, resulting largely from the varying patterns of breast tissue and their ability to disguise a cancer. Systems that can rapidly scan and analyze many mammograms can help radiologists reduce such errors. However, to date, such systems have not been practical for clinical use because they have neither achieved sufficient performance (sensitivity and specificity) nor the required processing speed for analyzing mammograms in a cost-effective, near real-time manner.

15

Since no two breast lesions look the same, computer-aided diagnosis of mammograms must be capable of generalizing to make correct decisions on data patterns of lesions that the computer has never before experienced. This is analogous in many ways to locating camouflaged targets in that neither "target" is well defined, and detection may require looking for large scale features and distortions of the overall scene (or total breast).

20

25 The prior art in the field includes the following:

Doi et al. "Computer-aided Diagnosis: Development of Automated Schemes for Quantitative Analysis of Radiographic Images," *Seminars in Ultrasound, CT, and MRI* 1992, 13:140-152.

30

Nishikawa et al., "Performance of Automated CAD Schemes for the Detection and Classification of Clustered Microcalcifications," In: *Digital Mammography*, Elsevier Science B.V., The Netherlands, 1994, 13-20.

35

Davies et al., "Automatic Computer Detection of Clustered Calcifications in Digital Mammograms," *Phy. Med. Biol.* 1990, 35:8:1111-1118.

Ng et al., "Automated Detection and Classification of Breast Tumors," *Computers and Biomedical Research* 1992, 25:218-237.

5 Kegelmeyer, Jr. et al., "Dense Feature Maps for Detection of Calcifications," In: *Digital Mammography*, Elsevier Science B.V., The Netherlands, 1994, 3-12.

Brettle et al., "Automatic Microcalcification Localisation using Matched Fourier Filtering," *Digital Mammography*, Elsevier Science B.V., The Netherlands, 1994, pp.21-30.

10 The prior art includes a variety of systems to detect both microcalcifications and mass lesions. For microcalcification or mass classification, the prior art techniques include either the use of features extracted by human observers or computer-extracted features. The latter features for microcalcifications include shape analysis such as compactness,
15 Fourier descriptors of image boundaries, and average distance between calcifications applied to the extracted features. For mass classification methods, features used include spiculations or irregular masses that are identified by local radiating structures or analysis of gradient histograms generated by seed growing and local thresholding methods, morphologically-based features and image texture. These features are classified using
20 either neural networks or binary decision trees. Common to all of these approaches is limited testing, with the computationally-intensive nature of the process implied as one reason for less than full, comprehensive testing.

25 Researchers associated with the University of Chicago (or Arch Development Corp.) have been major contributors in the development of computer-aided diagnosis for mammogram analysis (e.g., Doi et al.). This group has developed an automatic scheme for detecting clustered microcalcifications and is evaluating a method for classification of clusters. One approach involves using a linear filter to improve the signal-to-noise ratio of microcalcifications. Signal-extraction criteria are imposed to distinguish true lesions from artifacts. The computer then indicates locations that may contain clusters
30 of microcalcifications on the film. These investigators were able to demonstrate that use of the program improved radiologists' accuracy in detecting clustered microcalcifications. In the detection computer-aided diagnosis scheme, 78 digitized mammograms (half with, half without clusters) were processed and achieved 87% sensitivity to detection of clusters.

Davies et al. have described a system for detection of clusters of calcifications using image analysis techniques. Their method was based on a local area thresholding process, and was also able to identify other breast structures relative to the clusters of calcification. The results of their feasibility trial were a 100% true positive classification (for clusters), with an 8% false positive rate and 0% false negatives.

A more complex automated model is described by Ng et al. for the detection and classification of stellate lesions and circumscribed lesions. The relative strength of this system is that both stellate and circumscribed lesions appear as circular, bright masses with fuzzy boundaries but they may be difficult to distinguish by an automated system. Their pilot trial showed a high detection rate with a low false positive rate.

Kegelmeyer et al. have reported promising results when applying algorithms using binary decision trees and a dense feature map approach in identifying spiculated lesions and calcifications. Kegelmeyer et al. have achieved 100% sensitivity and 82% specificity when merging edge information identifying spicules with local texture measures, thus eliminating false-positive detections.

Brettle et al. demonstrated a true positive success rate of 100% and false positive and false negative rates of 0% each when using matched Fourier filtering in the frequency domain of mammographic images for detecting micro-calcification clusters. This application only used 15 segments of images, seven containing microcalcifications and eight without microcalcifications.

Other efforts by the National Institutes of Health and the National Science Foundation include: (i) gray scale image processing for better presentation of the image to the radiologist; (ii) applications of solutions of large-scale constrained optimization problems; (iii) adapting eye dwell time as a cue; and (iv) spatial filters and signal extraction techniques for detection of microcalcifications; noise smoothing, edge enhancement, and structured background correction methods for detection of the mass boundary and characteristics such as size, density, edge sharpness, calcifications, shape, lobulation, and spiculation extracted from the mass and used for classification.

All of the aforementioned reports indicate some level of success, though most use single scale techniques, not multiscale which is essential to a high confidence result.

From the data presented, it is difficult to compare the various results since there is no common data set. It is also difficult to assess the complexity of each data set.

3. **Solution to the Problem.** The present approach is based on an innovative concept to detect patterns in medical images (e.g., mammograms) using a Fourier transform optical correlator for image analysis followed by a series of neural networks hosted on a digital computer to analyze the results. An optical processor allows higher-order bandwidth, multi-resolution, multi-orientation approaches for feature extraction and enhancement that are not feasible in real time on a digital computer. A hybrid optical/digital computer approach ensures sufficient processing power at a moderate cost to accommodate discriminating algorithms and yet analyze a mammogram in a matter of seconds.

SUMMARY OF THE INVENTION

15 This invention provides a system for automated detection of cancerous masses in mammograms. The mammogram is digitized and regions of interest (ROIs) are detected using Fourier analysis (e.g., by means of an optical correlator). The pixels in the ROI are averaged together to create a smaller array of super-pixels that are input into a first neural net. Context data is extracted from the mammogram for each ROI, such as size, location, ranking, brightness, density, and relative isolation from other ROIs. A second neural net receives the output values from the first neural net and the context data as inputs and generates an output score indicating whether the ROI contains a cancerous mass. The second neural net can also be provided with context data from another view of the same breast, the same view of the other breast, or a previous mammogram for the same patient.

25 These and other advantages, features, and objects of the present invention will be more readily understood in view of the following detailed description and the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

30 The present invention can be more readily understood in conjunction with the accompanying drawings, in which:

FIG. 1 is a simplified block diagram of the present invention.

FIG. 2 is a flow chart of the present invention.

FIG. 3 is a section of a mammogram showing a region of interest that is potentially cancerous.

FIG. 4 is a section of a mammogram corresponding to FIG. 3 that has been overlaid with a super-pixel grid.

FIG. 5 is a section of a mammogram corresponding to FIG. 3 after the region of interest has been reduced to super-pixels.

FIG. 6 is a section of a mammogram corresponding to FIG. 3 after the region of interest has been enhanced.

FIG. 7 is an example of a report consisting of two views in a mammogram with the suspicious regions of interest marked.

DETAILED DESCRIPTION OF THE INVENTION

FIG. 1 provides a simplified block diagram of the overall system. A corresponding flow chart of the process used in the present invention is illustrated in FIG.

2.

1. **Acquire digital image of mammogram.** The first step in the process shown in FIG. 2 is acquisition of a digital mammogram 20 for analysis. A wide variety of hardware 11 can be used for image acquisition. For example, a high-resolution scanner can be used to scan a conventional mammogram film, or a direct digital image can be acquired. However, it is important that certain minimum standards be met. It is anticipated that 70 micron, 12 bit (reduced to 8 bit) digitization will be required if the present system is used for the purpose of aiding a radiologist in evaluating mammograms, and that 35 micron, 12 bit (reduced to 8 bit) digitization will be required for automated screening of mammograms.

Mammograms are typically paired by views (mediolateral and craniocaudal). As illustrated in FIG. 2, raw digitized images 20 are received by the workstation. These images are raw, single byte images in the 35-70 micron resolution range. Either a fixed size is used for the image or information on the pixel width and height of the image is included. A header is added to the image that allows the system to know the size and

data type of the image. A database is used to track the image, its size, and orientation (left breast, right breast, craniocaudal, or mediolateral). The quality of imagery is also examined. This is primarily a check that the image is a mammogram (i.e., it has the features of a mammogram) and that it meets specifications indicative of a mammogram ready to be read by a radiologist.

Image processing hardware 12 reduces the scale of the image to a resolution of approximately 230 microns. This reduces the processing requirement while maintaining sufficient resolution to identify regions of interest (ROIs). A blurring/contrast reduction algorithm is used where the image is superimposed on the background, to reduce "edge effects" (i.e., erroneous high frequency signals created by a sudden drop-off or edge).

2. **Identify regions of interest.** Regions of interest are identified in the image by Fourier analysis, as shown at reference number 21 in FIG. 2. This is done by comparing four bandpassed images with the original image and looking for "peaks". Bandpass filtering can be done in the digital domain using conventional computer hardware and signal processing algorithms. However, an optical correlator-13 such as that shown in U.S. Patent No. 5,418,380 (Simon et al.) is used in the preferred embodiment of the present invention to decrease processing time and increase through-put. This approach selects bright spots of specific sizes and passes them on to the neural networks 14, 15 for analysis.

Four bandpass filters are used to detect feature sizes. These bandpass filters are as follows:

(a) 38 micron resolution - detects features approximately in 0.5 - 5 mm range. The peak detector eliminates anything smaller than about 2.5 mm.

(b) 76 micron resolution - detects feature sizes approximately in 2.5 mm - 1 cm range.

(c) 152 micron resolution - detects feature sizes approximately in 5 mm - 2 cm range.

(d) 305 micron resolution - detects feature sizes approximately in 1 cm - 4 cm range.

-7-

The 230-micron images are correlated with each of the bandpass filters. For each of these four bandpasses, the strongest peaks are detected, then the sizes of the peaks are estimated by measuring the number of pixels around the peak that are:

- (a) at least 30% of the peak brightness of the bandpassed image; and
- 5 (b) at least 70% of the brightness of the original image measured at the peak location; and
- (c) at least 20% of the maximum brightness of the original image.

10 Peaks that are too small or too large relative to the scale of the bandpass are immediately rejected. A more detailed size determination is made on the remaining peaks. A radial size is calculated by measuring the distance from the peak in several directions until the pixel brightness is:

- (a) less than $1/\exp(2)$ of the peak brightness; or
- (b) less than 70% of the brightness of the original image measured at the peak
- 15 location; or
- (c) ~~greater than 5% higher than the next pixel closer to the peak ("peak~~
overlap").

20 Peaks that are too small or too large are eliminated, as are peaks lower in actual brightness than the surrounding area ("dark spots"), elongated features (maximum length across significantly greater than minimum length) and peaks which are too close to other peaks. Peaks that survive the selection criteria are then used to select regions of interest (ROIs). FIG. 3 is an example of a section of a mammogram showing a region of interest that is potentially cancerous. These ROIs are extracted from the original high-resolution

25 image in squares approximately 2-3 times feature size as measured by the first criteria and scaled to fixed size (256 x 256 pixels). There are eight possible ROI resolutions, ranging from a minimum size of 9.7 mm (38 micron resolution) to a maximum size of 7.8 cm (305 micron resolution).

30 **3. Analyze regions of interest.** Each region of interest is analyzed by means of a neural network using a super-pixeled image as an input. This grid is a radial-polar grid

with the angles evenly spaced and constant radial increments. In the preferred embodiment, the grid consists of 10 equally-spaced radial and 32 equally-spaced angular bands. The pixels in each grid space are averaged together to create 320 "super-pixels" (reference numeral 24 in FIG. 2). The super-pixels, along with two inputs determined from the approximate size of the feature, are used as the inputs to the first neural net 14. For example, FIG. 4 is a section of a mammogram corresponding to FIG. 3 that has been overlaid with a super-pixel grid. FIG. 5 is a section of a mammogram corresponding to FIG. 3 after the region of interest has been reduced to super-pixels.

The first neural network 14 is trained on both cancerous and non-cancerous regions of interest. Regions of Interest (ROIs) are selected from known (verified normal or verified cancerous) mammograms. These ROIs are ranked according to their appearance, with +1 being most normal-like and -1 being most lesion-like. The ROIs are then converted to super-pixels for input to the first neural network. A random scaling, rotation and translation of the super-pixel grid is performed, and half of the image (angularly) is chosen in a way that excludes any parts of the ROI outside the original area of the mammogram (These edges are generated by superimposing an odd-sized image on a fixed-sized background). The values of the half-grid are normalized in two steps: A stretching of the values over a fixed range increases the contrast within the half-grid, while a normalization of the sum of all the half-grid inputs assures consistency in brightness over all possible half-grids. A statistical thresholding of the ROI is used to determine the approximate "size" of the main feature in the ROI. This size measurement is converted to an equivalent diameter, perturbed by the same scaling perturbation used on the super-pixel grid, then converted into a pair of inputs, using a sine-cosine transformation. These two values, along with the grid super-pixels, are input to the first neural network using a randomly initialized set of neural network weights. The output of the first neural network (restricted between +1 and -1) is compared with the chosen ranking. An error value is calculated and used to correct the neural network weights using a back-propagation learning algorithm known as REM (Recursive Error Minimization). The REM algorithm uses the derivatives of the propagated error to create a more stable convergence than conventional back-propagation. A new set of perturbations is then applied to the next ROI, creating a new half-grid, which is again

normalized and applied to the network. This process is repeated until the RMS of the error values drops to a small value, or the change in the error over several iterations becomes very small, whichever occurs first. This signifies that the network is now trained. The number of iterations depends on several conditions, but in general will take about 1000-3000 cycles through the complete set of ROIs.

Once trained, a set of ROIs selected from "unknown" mammograms is processed through the neural network (reference numeral 25 in FIG. 2). No perturbations are done on these inputs, although the grid is rotated through all possible angles, one angular grid space at a time, such that no half-grid includes any "edges". Each half-grid in turn is normalized as above, and applied to the trained network along with the two size inputs from the ROI. The outputs for the rotations are combined, producing two or three statistical outputs for each ROI. These outputs are then used as part of the set of inputs for the second neural network 15, as discussed below.

4. **Analyze regions of interest in context of mammogram.** A combination of 16 wavelet filters and two bandpass filters are applied to each region of interest to create enhanced ROIs at reference numeral 22 in FIG. 2. FIG. 6 is a section of a mammogram corresponding to FIG. 3 after the ROI has been enhanced to emphasize the contours and lines of the features in the image. At reference numeral 23 in FIG. 2, these enhanced images are used, along with the original ROIs, to generate context inputs for the second neural net 15 (or "context network"). The context inputs for the context network 15 includes the location, ranking, brightness and density profiles, sizes of inner and outer features, relative isolation from other ROIs, roughness of contours, fuzziness, and the number and length of any spicules. The distance of the center of each ROI from the nipple is determined by:

- (a) For the mediolateral view, locating the chest wall (if possible), and measuring the distance of the ROI from the nipple perpendicular to the chest wall;
- (b) For the craniocaudal view, measuring the horizontal distance of the ROI from the nipple.

For each ROI, up to five ROIs in the alternate view are selected that are within 50% of the same distance from the nipple. Sizes and output values from the first neural network 14 are also used as inputs to the context neural net 15. The context inputs to the second neural net 15 are calculated by using sine and cosine transformations of the parameters defined above.

The context neural net 15 is initially trained, again using verified cancerous and non-cancerous cases with random deviations for each input to generate training examples. The context inputs for each ROI are determined, including the outputs from the trained first neural network 14. Each input is perturbed by a random value (e.g., +/-10%) and then converted into a pair of inputs using a sine-cosine transformation. The full set of inputs is applied to the second neural network 15 using a randomly initialized set of neural network weights. The output of the second neural network 15 is compared to the same ranking value used for the first neural network 14. The weights are corrected by the method cited above, and the second network 15 is trained to convergence using about 1000-3000 cycles through the set of ROIs.

The "unknown" ROIs are passed once through the trained second network 15 without perturbations. The output of this network is then combined with that of other ROIs belonging to the same mammogram (including other views, other breast and/or other years) to characterize the ROI and the mammogram. After training is complete, the trained context neural network 15 can then be used to evaluate unknown cases (reference numeral 26 in FIG. 2).

In the preferred embodiment of the present invention, the context neural net 15 is a second network separate from the first neural network 14 used in step 3. However, it should be expressly understood that other configurations could be substituted. For example, a single neural net could receive both the ROI super-pixels and context data as inputs.

5. Threshold regions of interest. After training is complete, an ROC curve is generated from the output score produced by the second neural network in response to sets of known mammograms. An appropriate threshold output score can be determined

-11-

from the resulting ROC curve which will achieve the desired probability of detection and false alarm rate.

This threshold can then be used to classify unknown cases during normal operation of the system (reference numeral 27 in FIG. 2). For example, the threshold score is used by the decision software 16, as depicted in FIG. 1, to generate a report for the radiologist.

6. **Create report.** A user-friendly graphical report is provided to the radiologist on a CRT or other display 17 (reference numeral 28 in FIG. 2). FIG. 7 is an example of a report consisting of two views in a mammogram with the suspicious regions of interest marked. This report typically shows both views of the mammogram, along with a classification determined by the analysis of the neural networks. The locations of suspect masses are marked on the display. For example, the final classifications might include the following:

15 (a) "Suspect" mammogram has at least one ROI detected on either view with a value less than a predetermined threshold (i.e., more lesion-like), but more than a fixed level (e.g., 0.2) below the threshold. A mammogram with such a designation needs further review by the radiologist.

20 (b) "Cancer" mammogram has at least one ROI detected on either view with a value less than the minimum defined above for a "suspect." A mammogram with such a designation needs further review by the radiologist.

(c) "Normal" mammogram has no ROIs on either view with values less than the threshold. This mammogram needs no further review.

25 7. **Results.** A round-robin test was run with over 100 different pairs (two views per breast) of mammograms analyzed. The results of this test showed that, when the threshold was set to detect 100% of the cancers, approximately 50% of the mammograms were screened, with less than 0.7 false positives per image.

30 **Application.** The present system can be used by a radiologist in either of two modes of operation. In the first mode, the system merely aides the radiologist by

flagging any ROIs falling into the "Suspect" or "Cancer" classifications. The radiologist continues to thoroughly review each mammogram and remains responsible for diagnosis. The system is equipped with an image archive 18 shown in FIG. 1. The archive can store each mammogram or only those mammograms designated by the radiologist. In the second mode ("screening mode"), the system analyzes each mammogram and only those classified as "Suspect" or "Cancer" are brought to the attention of the radiologist for further review. "Normal" mammograms are sent directly to the image archive 18 for archival storage.

Other Context Inputs. In addition to the embodiment described above, a number of additional types of context data can be input into the context neural net to enhance the accuracy of the system. For example, the ROIs can be analyzed in the context of information from the other view of same breast. This step can be combined with the previous step using the second neural network 15 as described above, or performed separately using a third neural network. The additional context data can include information on the size and brightness of the ROI under question as well as other ROIs in the same breast and ROIs in the alternate view which have nearly the same perpendicular distance to the nipple. For training purposes, each value is allowed to vary over a specified range to gain additional test cases. The second neural network 15 is then trained using both cancerous and non-cancerous ROIs. The resulting connections are then used to evaluate ROIs outside of the training set.

Similarly, the ROIs can be analyzed in the context of the same view of other breast for the same patient. This step is implemented by performing the same operations on the other breast, then cross-correlating the information to determine asymmetries between the two breasts. For example, additional context inputs can be input into the context neural net 15 corresponding to the sizes and density network results for ROIs in similar locations on the other breast. These additions should produce a significant reduction in the false alarm rate, since it is very unlikely that cancers would symmetrically develop in both breasts. This implementation is conceptually identical to the use of context inputs derived from alternative views of the same breast as previously

-13-

discussed, but requires additional sets of mammograms for adequate training of the context neural net 15.

5 The ROIs can also be analyzed in the context of prior year mammograms for the same patient. Additional context inputs can be input into the context neural net 15 corresponding to ROIs in similar locations in prior-year mammograms for the same patient. This is similar to the use of context inputs derived from contemporaneous alternate views, as described above, but instead relies on detecting changes that have occurred since previous mammograms. These additional context inputs should produce a significant reduction in the false alarm rate in light of the temporal development
10 properties of cancers in contrast to non-cancers. Here again this implementation requires a additional sets of mammograms for adequate training of the context neural net 15.

The previous discussion has been directed primarily to detection of cancerous masses in mammograms. It should be expressly understood that the present invention could be employed to detect abnormalities in a wide variety of medical images. For
15 example, the present system could be used to detect cancer or other types of lesions in a chest x-ray film; or cancerous and pre-cancerous cells in a pap smear or biopsy.

The above disclosure sets forth a number of embodiments of the present invention. Other arrangements or embodiments, not precisely set forth, could be practiced under the teachings of the present invention and as set forth in the following
20 claims.

WE CLAIM:

1. A method for automated analysis of medical images to detect an abnormality comprising:
 - detecting a region of interest (ROI) containing a possible abnormality in a medical
5 image using Fourier analysis;
 - extracting context data for said ROI from said medical image; and
 - inputting said ROI and said context data to a neural net and generating an output indicating whether said abnormality is present in said ROI.
- 10 2. The method of claim 1 wherein said medical image comprises a mammogram and the abnormality is the presence of a cancerous mass.
3. The method of claim 1 wherein the step of detecting a ROI is performed by an optical correlator.
- 15 4. The method of claim 1 further comprising the step of averaging pixels in said ROI to create a smaller number of super-pixels that are input to said neural net.
5. The method of claim 4 wherein said super-pixels are arranged in a radial-polar
20 pattern.
6. The method of claim 1 wherein said context data comprises the location of said ROI relative to a standard anatomical landmark in said medical image.
- 25 7. The method of claim 1 wherein said context data comprises the size of said ROI.
8. The method of claim 1 wherein said context data comprises the brightness and density profiles of said ROI.
- 30 9. The method of claim 1 wherein said context data comprises the relative isolation of said ROI from other ROIs in said medical image.

-15-

10. The method of claim 1 wherein said context data comprises the number and length of any spicules in said ROI.

5 11. The method of claim 1 wherein said context data comprises information about other ROIs in said medical image.

12. The method of claim 1 further comprising the step of enhancing the contours of said ROI before extracting said context data.

10 13. The method of claim 1 wherein said neural net comprises:
a first neural net receiving said ROI as an input and generating at least one output value; and
a second neural net receiving said output values from said first neural net and said context data as inputs and generating an output indicating whether said abnormality is
15 present in said ROI.

14. The method of claim 1 wherein said neural net is initially trained using perturbations of a set of training images.

20 15. A system for automated detection of cancerous masses in mammograms comprising:

means for inputting a digital mammogram;
an optical correlator to detect any region of interest (ROI) in said mammogram;
means for extracting context data for said ROI from said mammogram; and
25 a neural net receiving said ROI and said context data as inputs and generating an output indicating whether said ROI contains a cancerous mass.

16. The system of claim 15 further comprising means for averaging pixels in said ROI to create a smaller number of super-pixels that are input to said neural net.

-16-

17. The system of claim 16 wherein said super-pixels are arranged in a radial-polar pattern.
18. The system of claim 15 wherein said context data comprises the location of said ROI relative to the nipple in said mammogram.
19. The system of claim 15 wherein said context data comprises the size of said ROI.
20. The system of claim 15 wherein said context data comprises the brightness and density profiles of said ROI.
21. The system of claim 15 wherein said context data comprises the relative isolation of said ROI from other ROIs in said medical.
22. The system of claim 15 wherein said context data comprises the number and length of any spicules in said ROI.
23. The system of claim 15 wherein said context data comprises information about other ROIs in said mammogram.
24. The system of claim 15 wherein said mammogram includes a plurality of views of the same breast, and wherein said context data comprises information from another view of the same breast.
25. The system of claim 15 wherein said mammogram includes views of both breasts, and wherein said context data comprises information about the other breast.
26. The system of claim 15 wherein said context data includes information from a previous mammogram for the same patient.

-17-

27. The system of claim 15 further comprising filter means for enhancing the contours of said ROI before extracting said context data.

28. The system of claim 15 wherein said neural net comprises:

5 a first neural net receiving said ROI as an input and generating at least one output value; and

a second neural net receiving said output values from said first neural net and said context data as inputs and generating an output indicating whether said ROI contains a cancerous mass.

10

29. The system of claim 15 wherein said neural net is initially trained using perturbations of a set of training images.

30. A system for automated detection of cancerous masses in mammograms comprising:

15

means for inputting a digital mammogram containing an array of pixels;

an optical correlator to detect any region of interest (ROI) in said mammogram;

means for extracting context data for said ROI from said mammogram;

means for averaging pixels in said ROI to create a smaller array of super-pixels;

20

a first neural net receiving said super-pixels as an input and generating at least one output value; and

a second neural net receiving said output values from said first neural net and said context data as inputs and generating an output indicating whether said ROI contains a cancerous mass.

25

31. The system of claim 30 wherein said super-pixels are arranged in a radial-polar pattern.

32. The system of claim 30 wherein said context data comprises information about other ROIs in said mammogram.

30

33. The system of claim 30 wherein said mammogram includes a plurality of views of the same breast, and wherein said context data comprises information from another view of the same breast.

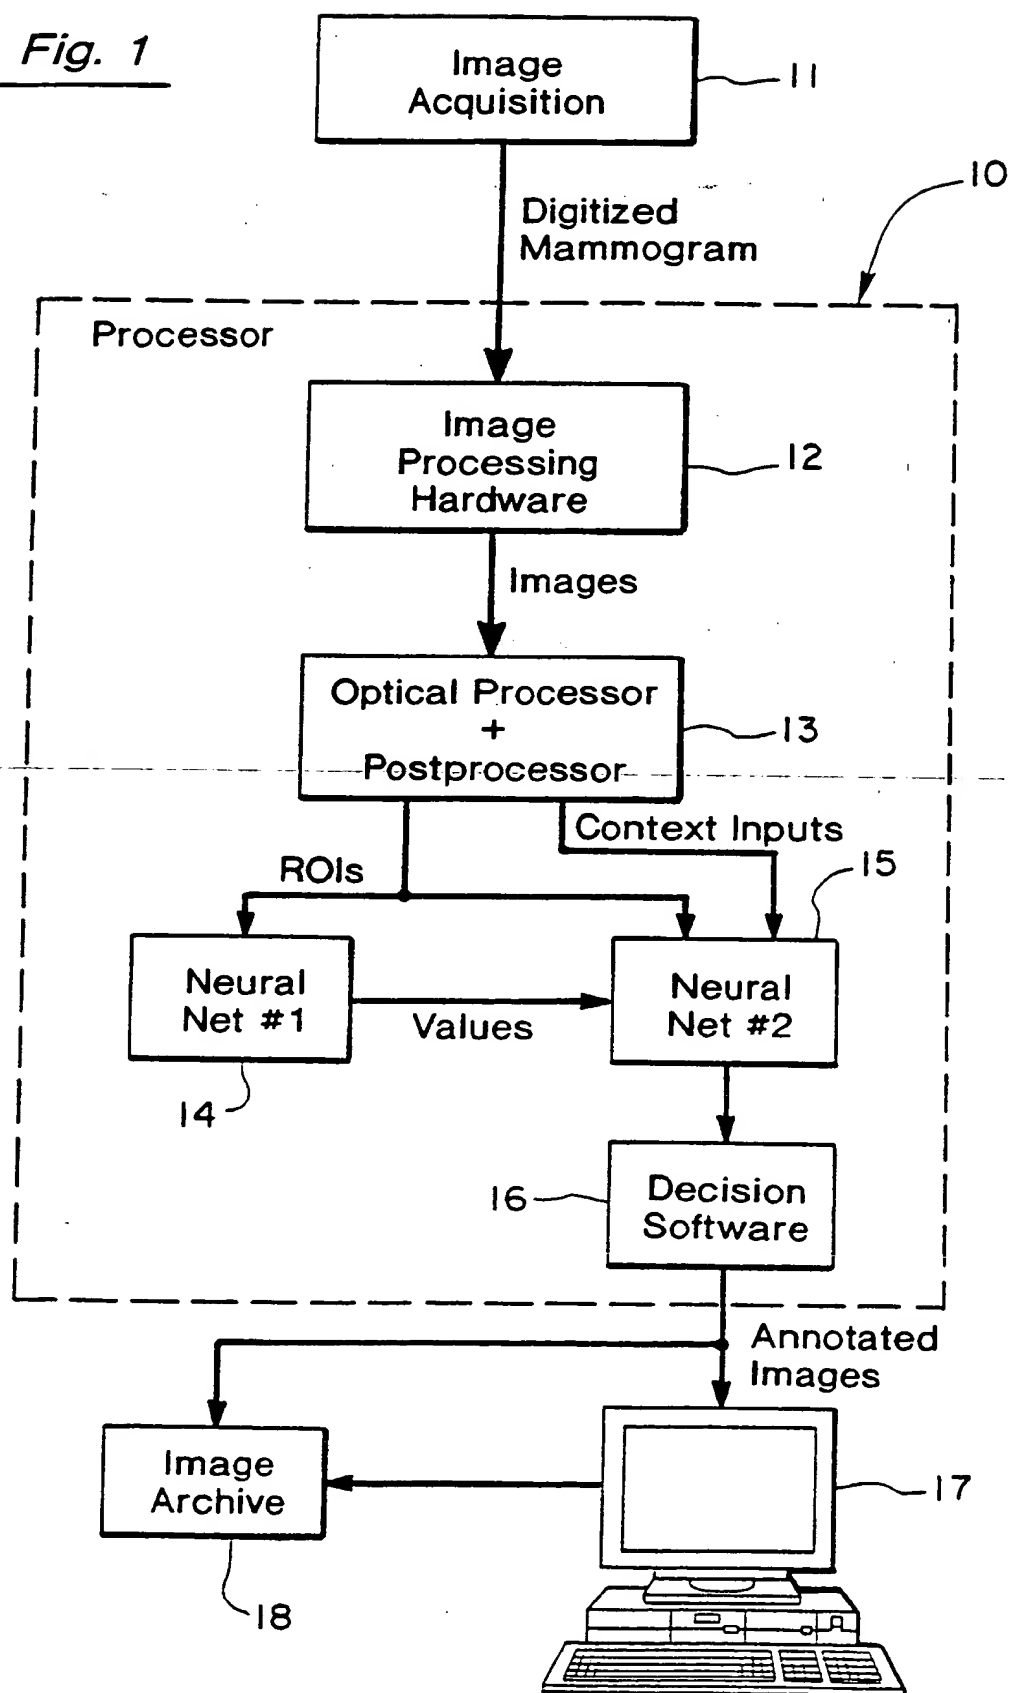
5 34. The system of claim 30 wherein said mammogram includes views of both breasts, and wherein said context data comprises information about the other breast.

35. The system of claim 30 wherein said context data includes information from a previous mammogram for the same patient.

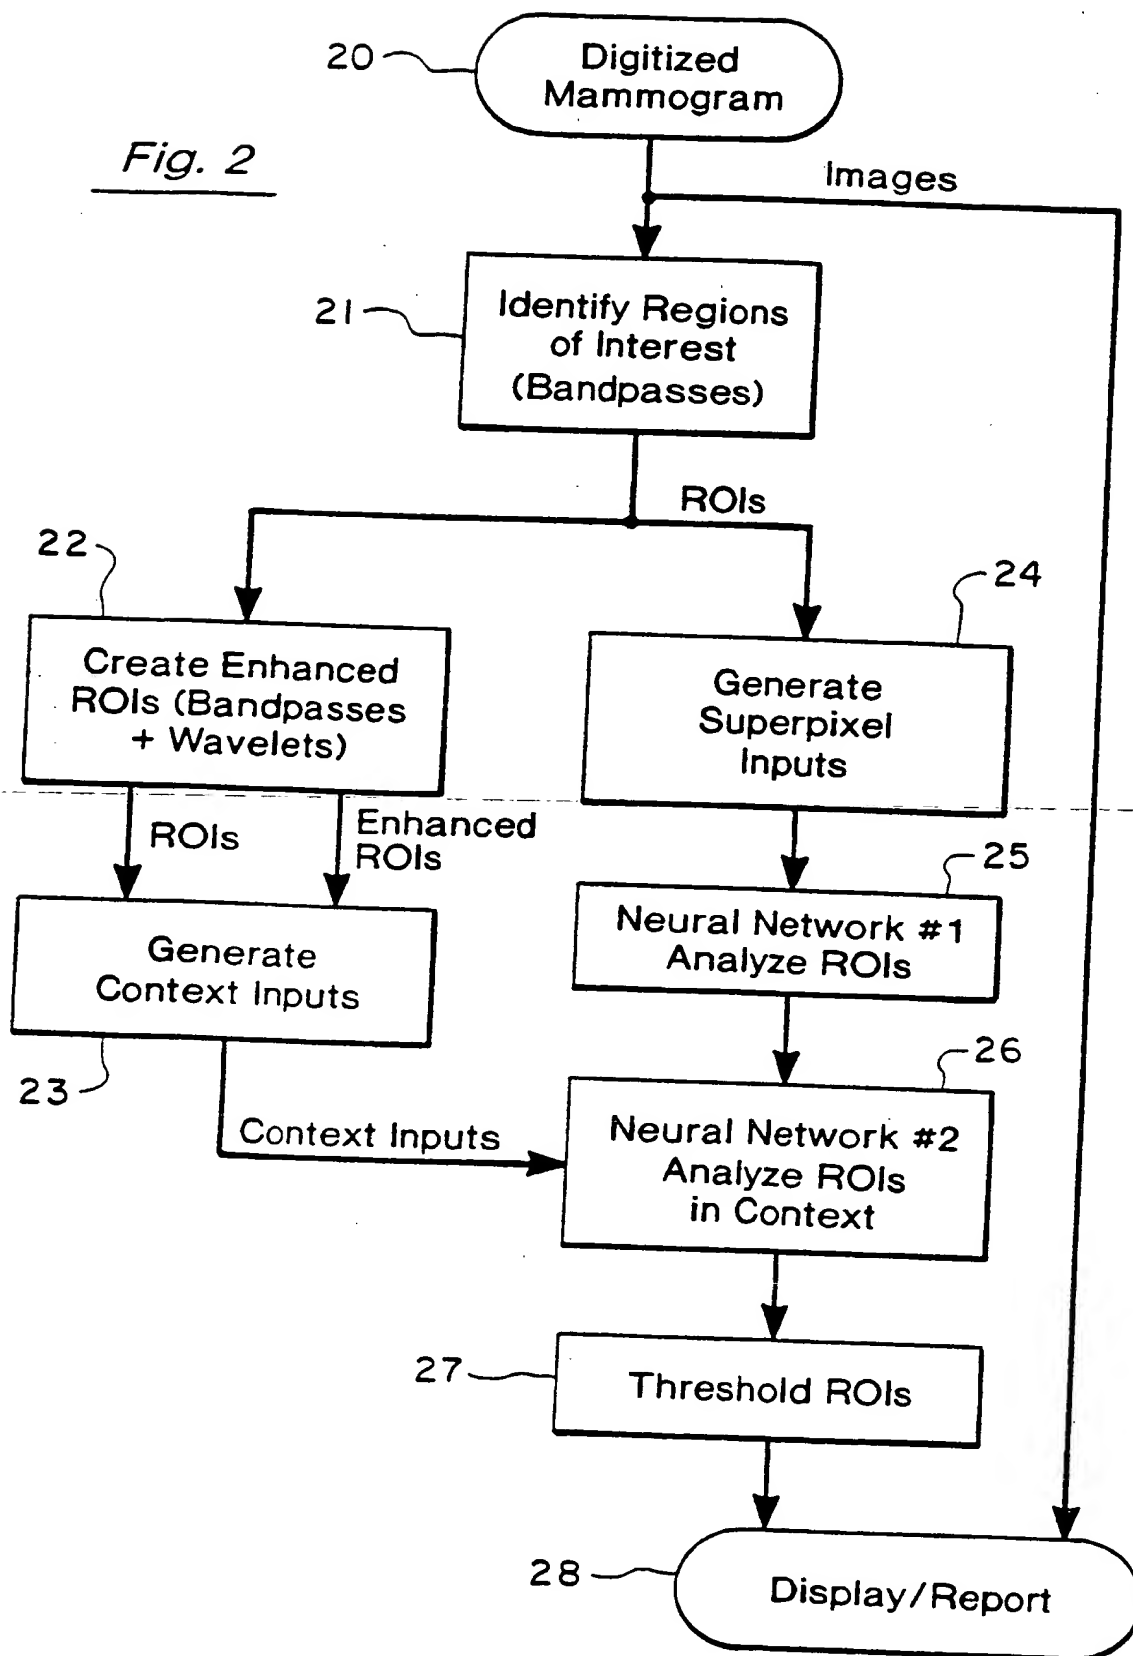
10

36. The system of claim 30 further comprising filter means for enhancing the contours of said ROI before extracting said context data.

1/5

Fig. 1

2/5

Fig. 2

3/5

Fig. 3





Fig. 4



4/5

Fig. 5





Fig. 6



5/5




Fig. 7

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/15350

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06K 9/00

US CL : 382/128, 129, 130, 131, 132, 133;

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 382/128, 129, 130, 131, 132, 133;

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
none

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS.

Search Terms: mammogram, ROI, pixel, neural net, Fourier

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| Y | US 5,491,627 A (ZHANG et al.) 13 February 1996, fig. 1, elements 10-16; figure 3; col. 2, lines 20-64; col. 10, lines 1-64; col. 17, lines 1-40; | 1-35 |
| Y | US 5,754,693 A (TAKESUE et al.) 19 May 1998, col. 1, lines 15-40; col. 4, lines 10-14; col. 15, lines 5-40; col. 16, lines 20-35; | 1-35 |
| Y | US 5,003,979 A (MERICKEL et al.) 02 April 1991, fig. 8; fig. 12(A); fig. 12(B); col. 5, lines 35-60; col. 8, lines 15-35; | 1-35 |
| Y | US 5,212,637 A (SAXENA) 18 May 1993, col. 6, lines 44-66; col. 11, lines 10-40; | 1, 15, 30 |

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

| | | |
|---|-----|--|
| * Special categories of cited documents: | *T | later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention |
| *A* document defining the general state of the art which is not considered to be of particular relevance | *X* | document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone |
| *E* earlier document published on or after the international filing date | *Y* | document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | *A* | document member of the same patent family |
| *O* document referring to an oral disclosure, use, exhibition or other means | | |
| *P* document published prior to the international filing date but later than the priority date claimed | | |

Date of the actual completion of the international search

09 OCTOBER 1998

Date of mailing of the international search report

13 APR 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JAY PATEL

James R. Matthews

Telephone No. (703) 308-7728

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/15350

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

[illegible]